arXiv:1610.03793v1 [cs.LG] 12 Oct 2016

Introduction to the "Industrial Benchmark"

Daniel Hein, Alexander Hentschel, Volkmar Sterzing, Michel Tokic, Steffen Udluft

Siemens AG, CT RTC BAM LSY-DE
Learning Systems

October 13, 2016

Abstract

A novel reinforcement learning benchmark, called Industrial Benchmark, is introduced. The Industrial Benchmark aims at being be realistic in the sense, that it includes a variety of aspects that we found to be vital in industrial applications. It is not designed to be an approximation of any real system, but to pose the same hardness and complexity.

1 Introduction

The scientific method requires that a hypothesis is tested by experiments. This holds true in the field of machine learning, when algorithms are to be developed or improved. Such a test can be, to run the algorithm on a real system, in order to observe its performance, or to run it on a simulation, a virtual system implemented as a computer program. The latter method has several advantages: it is usually faster, cheaper, and of course more safe to test the algorithm on a simulation. In addition the simulation can be manipulated much more freely than a real system. Internal states can be observed, stored, restored, and set, thus given the freedom to test special aspects with little effort. The amount of data can usually be enlarged such that all results gain statistical significance. This stands in drastic contrast to the situation when testing the algorithm on a large scale industrial system, like a power plant.

There are some disadvantages with simulation based testing though. First of all, the final success of an improved algorithm will be defined by the performance on the real system. Thus any deviation of the simulation from the real system might cause the development of new algorithms and the process of improving to go not in the right direction. The simulation might be too simple, underestimating the challenges posed by the real system. Or, it might focus on less relevant aspects of the task, posing artificial difficulties. Therefore it would be desirable to use a most realistic simulation for benchmarking. On the other hand, when developing an algorithm one usually targets for methods that are applicable for a broad spectrum of systems and fine tuning for a specific system is fruitful from the perspective of that system only, while it does not help to decide for the algorithm which is "superior in general". It can be doubted that it is possible to create an algorithm that is "superior in general". Usually special cases exist where the seemingly inferior method is superior, but still it seems good practice to aim for a method that performs well in a wide variety of simulation benchmarks. This dilemma of specific versus general solutions will not be solved here, but we want to contribute a software benchmark[1] that

[1] Java source code: http://github.com/siemens/industrialbenchmark

captures many aspects that we found to be vital in industrial applications. The basic task, we are considering is the optimization of operating an industrial system. This task will be described in the theoretical framework of reinforcement learning [1]. The proposed simulation, called Industrial Benchmark, will serve as environment, i.e. the system to be controlled. Independently of reinforcement learning, the Industrial Benchmark can also be used to evaluate regression, forecasting, and system identification capabilities of different machine learning methods, as well as specific challenges like transfer learning, active learning, feature selection, or change detection.

2 Industrial Benchmark

The Industrial Benchmark aims at being realistic in the sense, that it includes a variety of aspects that we found to be vital in industrial applications. It is not designed to be an approximation of any real system, but to pose the same hardness and complexity. State- and action-space are continuous, the state-space is rather high-dimensional, and only partially observable. The actions consist of three continuous components and effect three steerings. There are delayed effects. The optimization task is multi-criterial in the sense that there are two reward-components that show opposite dependencies on the actions. The dynamical behaviour is heteroscedastic with state dependent observation noise and state dependent probability distributions, based on latent variables. The dynamical behaviour is dependent on an external driver, that cannot be influenced by the actions. The Industrial Benchmark is designed such that the optimal policy will not approach a fixed operation point in the three steerings. Any specific choice is driven by our experience with industrial challenges.

3 Detailed description

At any time step t the reinforcement learning agent can influence the environment (Industrial Benchmark) via actions $\vec{a}(t)$ that are three dimensional vectors in $[-1,1]^3$. Each action can be interpreted as three proposed changes to the three observable state variables called current steerings. Those current steerings are named velocity v, gain g, and shift s. Each of those is limited to $[0,100]$.

$$\vec{a}(t) = (\Delta v, \Delta g, \Delta s)^\top , \tag{1}$$

$$\begin{aligned} v(t+1) &= \min(0.\max(100, v + \Delta v)) , && (2)\\ g(t+1) &= \min(0.\max(100, g + 10\,\Delta g)) , && (3)\\ s(t+1) &= \min(0.\max(100, s + \alpha_s \Delta s)) . && (4) \end{aligned}$$

Where the step size for changing shift is $\alpha_s = 20 \sin(15^0)/0.9 \approx 5.75$. After applying the action $\vec{a}(t)$ the environment transitions to the next time step $t+1$ in which it is in an internal state $\vec{s}(t+1)$[2]. State $\vec{s}(t)$ and successor state $\vec{s}(t+1)$ (also written as $\vec{s'}(t)$) are the markovian states of the environment that are only partially observable to the agent.

Three observable state variables we have already discussed, the current steerings v, gain g, and shift s. There are three more observable state variables. One is the external driver named set point p, that influences the dynamical behaviour of the environment but cannot be changed by the actions. In the setting discussed here, the set point is kept constant (Constant Set Point Setting).

[2]Not to be mixed up with the shift s, which is one of the variables that build the state $\vec{s}$.

An extension will be that the set point will change with time (Variable Set Point Setting), i.e. $p = p(t)$. In this Variable Set Point Setting the changes in the set point will be influenced externally and not by the actions. In the Constant Set Point Setting the agent does not need to predict set point changes, as they do not occur. This setting is therefore not one partially observable markov decision problem (POMDP), but a family (german: Schar) of POMDPs, parameterized by the set point. Learning to act optimal for any set point in the Constant Set Point Setting can also be seen as a multi-task or transfer learning task [2, 3].

The set of observable state variables is completed by the two reward relevant variables consumption $c(t)$ and fatigue $f(t)$. In the general reinforcement learning setting a reward $r(t)$ is drawn for each transition $t \to t+1$ from state $\vec{s}(t)$ via action $\vec{a}(t)$ to the successor state $\vec{s}(t+1)$ from a probability distribution depending on $\vec{s}(t)$, $\vec{a}(t)$, and $\vec{S}(t+1)$. In the Industrial Benchmark the reward is given by a deterministic function of the successor state $r(t) = r(\vec{s}(t+1))$, i.e.

$$r(t) = -(c(t+1) + f(t+1))\,. \tag{5}$$

In the real world tasks that motivated the Industrial Benchmark the reward function has always been known explicitly. In some cases it was subject to optimization itself and had to be adjusted to properly express the optimization goal. For the Industrial Benchmark we therefore assume that the reward function is known and all variables influencing it are observable.

Thus the observation vector $\vec{O}(t)$ at time t comprises current values of the set of observable state variables, which is a subset of all the variables of state $\vec{s}(t)$, i.e.

1. the current steerings, velocity $v(t)$, gain $g(t)$, and shift $s(t)$,
2. the external driver, set point p,
3. and the reward relevant variables consumption $c(t)$ and fatigue $f(t)$.

The data base for learning comprises of tuples $(\vec{O}(t),\vec{a}(t),\vec{O}(t+1),r(t))$, which, by introducing the notation $\vec{O}'$ for the observation vector of the successor state, will be written as $(\vec{O}(i),\vec{a}(i),\vec{O}'(i),r(i))$ or, in short, $(\vec{O},\vec{a},\vec{O}',r)(i)$.

The agent is allowed to use all previous observation vectors and actions to estimate the markovian state $\vec{s}(t)$.

4 Description of the dynamical behaviour

The dynamical behaviour of the Industrial Benchmark is determined by the three steerings velocity v, gain g, and shift s, the external driver set point p, and five latent variables. The dynamics can be decomposed in three different sub-dynamics named operational costs, mis-calibration, and fatigue.

4.1 Dynamics of operational cost

The sub-dynamics of operational cost is influenced by the external driver set point p and two of the three steerings, namely velocity v and gain g. The current operational cost $o(t)$ is calculated as

$$o(t) = \exp\left(\frac{2p(t) + 4v(t) + 2.5g(t)}{100}\right)\,. \tag{6}$$

The current operational cost $o(t)$ cannot be observed, the observation is delayed and smeared out by a convolution

$$\begin{aligned} o^c(t) \;=\; & 0o(t) + 0o(t-1) + 0o(t-2) + 0o(t-3) + 0o(t-4) + \\ & \frac{1}{9}o(t-5) + \frac{2}{9}o(t-6) + \frac{3}{9}o(t-7) + \frac{2}{9}o(t-8) + \frac{1}{9}o(t-9) \end{aligned} \tag{7}$$

The convoluted operational cost $o^c(t)$ still cannot be observed directly, it is modified by the second sub-dynamics, called mis-calibration, and finally subject to observation noise. The motivation for this dynamical bahaviour is that it is non-linear, depends on more than one influence, is delayed and smeared. All those effects have been observed in industrial applications.

4.2 Dynamics of mis-calibration

The sub-dynamics of mis-calibration is influenced by the external driver, set point p, and only one steering, namely shift s. Set point p and shift s are combined to an effective shift s^e

$$s^e = \min(1.5, \max(-1.5, s/20 - p/50 - 1.5)) \,, \tag{8}$$

which influences the three latent variables m_1^l, m_2^l, and m_3^l. The resulting mis-calibration m is a function of effective shift and the latent variables

$$m = \mathrm{f}(s^e, m_1^l, m_2^l, m_3^l) \,. \tag{9}$$

The resulting mis-calibration $m(t)$ is added to the convoluted operational cost $o^c(t)$, giving $\hat{c}$,

$$\hat{c} = o^c(t) + 25 m(t) \,, \tag{10}$$

Before being observable as consumption c the modified operational cost $\hat{c}$ is subject to heteroskedastic observation noise

$$c = \hat{c} + \mathrm{gauss}(0, 1 + 0.02\,\hat{\mathrm{c}}) \,, \tag{11}$$

i.e. a gaussian noise with zero mean and a standard deviation of $\sigma = 1 + 0.02\,\hat{c}$.

4.3 Dynamics of fatigue

The sub-dynamic of fatigue is influenced by the same variables as the sub-dynamic of operational cost, i.e. set point p, velocity v, and gain g. The Industrial Benchmark is designed in such a way that when changing the steerings velocity v and gain g as to reduce the operational cost, fatigue will be increased, leading to the desired multi-criterial task, with two reward-components showing opposite dependencies on the actions. The basic fatigue f_b is computed as

$$f_b = \max\left(0, \frac{30000}{5\,v + 100} - 0.01\,g^2\right) \,. \tag{12}$$

From the basic fatigue f_b, the fatigue f is calculated by

$$f = f_b(1 + 2\alpha) \,, \tag{13}$$

where α is an amplification. The amplification depends on two latent variables h_v and h_g, an effective velocity v^e, an effective gain g^e, and is affected by noise,

$$\alpha = \begin{cases} \frac{1}{1+\exp(-\mathrm{gauss}(2.4,0.4))} & \text{for} \quad \max(h_v, h_g) = 1.2 \\ \max(\eta^v, \eta^g) & \text{else} \,. \end{cases} \tag{14}$$

The noise components η^v and η^g, as well as the latent variables h_v and h_g, depend on effective velocity v^e, and effective gain g^e. These are calculated by set point dependent transformation functions

$$\mathrm{T}_v(v, g, p) = \frac{g + p + 2}{v - p + 101} \,, \tag{15}$$

$$\mathrm{T}_g(g, p) = \frac{1}{g + p + 1} \,. \tag{16}$$

Based on this transformation functions, effective velocity v^e and effective gain g^e are computet as follows:

$$v^e = \frac{\mathrm{T}_v(v,g,p) - \mathrm{T}_v(0,100,p)}{\mathrm{T}_v(100,0,p) - \mathrm{T}_v(0,100,p)} \tag{17}$$

$$g^e = \frac{\mathrm{T}_g(g,p) - \mathrm{T}_g(100,p)}{\mathrm{T}_g(0,p) - \mathrm{T}_g(100,p)}\,. \tag{18}$$

To compute the noise components η^v and η^g, six random numbers are drawn from different random distributions: η_e^v and η_e^g are drawn from an exponential distribution with mean 0.05, η_b^v and η_b^g are drawn from binominial distributions $\mathrm{Binom}(1, v^e)$ and $\mathrm{Binom}(1, g^e)$, respectively, η_u^v and η_u^g are drawn from an uniform distribution in $[0,1]$. These random numbers are combined to two noise components η^v and η^g by

$$\eta^v = \eta_e^v + (1-\eta_e^v)\eta_u^v\eta_b^v v^e \tag{19}$$

$$\eta^g = \eta_e^g + (1-\eta_e^g)\eta_u^g\eta_b^g g^e\,. \tag{20}$$

The latent variables h_v and h_g are caclulated as

$$h_v(t) = \begin{cases} v^e & \text{for} \quad v^e \le 0.05 \\ \min(5, 1.1h_v(t-1)) & \text{for} \quad v^e > 0.05 \wedge h_v(t-1) \ge 1.2 \\ 0.9h_v(t-1) + \frac{\eta^v}{3} & \text{else} \end{cases} \tag{21}$$

$$h_g(t) = \begin{cases} g^e & \text{for} \quad g^e \le 0.05 \\ \min(5, 1.1h_g(t-1)) & \text{for} \quad g^e > 0.05 \wedge h_g(t-1) \ge 1.2 \\ 0.9h_g(t-1) + \frac{\eta^g}{3} & \text{else}\,. \end{cases} \tag{22}$$

The sub-dynamic of fatigue results in a value for fatigue f, which is relevant for the reward function. (see Eq. 5).

5 State definitions

To give an overview on possible state definitions a small summary is given.

5.1 The observation vector

Only a part of the state variables is observable. This observation vector is also called observable state, but one has to keep in mind, that it does not fulfill the markov property. The observation vector $\vec{O}(t)$ at time t comprises current values of velocity $v(t)$, gain $g(t)$, shift $s(t)$, set point $p(t)$, consumption $c(t)$, and fatigue $f(t)$.

5.2 The preferred minimal markovian state

The preferred minimal markovian state fulfills the markov property with the minimum number of variables. It comprises 20 values. These are the observation vector (velocity $v(t)$, gain $g(t)$, shift $s(t)$, set point $p(t)$, consumption $c(t)$, and fatigue $f(t)$) plus some latent variables of the sub-dynamics. The sub-dynamics of operational cost adds a list of previous operational costs, $o(t-i)$ with $i \in 1, \cdots, 9$. Note that the current operational cost $o(t)$ is not part of this state definition. It would be redundant, as it can be calculated by $v(t)$, gain $g(t)$, and set point p. The sub-dynamics of mis-calibration needs 3 additional latent variables, m_1, m_2, and m_3, (Sec. 4.2). The sub-dynamics of fatigue adds 2 additional latent variables h_v and h_g, (Eq. 21 and 22).

text name	symbol	software name
set point	$p(t)$	`SetPoint`
velocity	$v(t)$	`Velocity`
gain	$g(t)$	`Gain`
shift	$s(t)$	`Shift`
consumption	$c(t)$	`Consumption`
fatigue	$f(t)$	`Fatigue`

Table 1: The observation vector.

5.3 The extended state

The extended state, also called the internal markovian state, contains in addition to all the variables of the preferred minimal markovian state also some variables, which are of useful for data analysis purposes.

6 Experimental setup

To test the algorithms in an initial batch mode, off-policy setting, data is generated by the maximum entropy policy,

$$\frac{dP(a|s)}{da} = \text{const} \tag{23}$$

The benchmark is initialized for ten different set points $p \in 10{,}20, \cdots ,90{,}100$ with the latent variables in their default values and the three steerings at 50 each. Then for each set point value the maximum entropy policy is applied on the benchmark for 1000 time steps, resulting in 10.000 data points. This data can then be used to train system identification models and/or policies. The goal is to build a policy π that maximizes the average reward on the same setting, where instead of the maximum entropy policy, the policy π is applied.

7 Interfaces

The main interfaces are defined in `com.siemens.rl.interfaces`. They are DataVector, Environment, and ExternalDriver.

7.1 Interface DataVector

This interface lists all necessary methods to implement a data vector, which might be a state- or action-vector.

text name or description	symbol	software name
set point	$p(t)$	`SetPoint`
velocity	$v(t)$	`Velocity`
gain	$g(t)$	`Gain`
shift	$s(t)$	`Shift`
consumption	$c(t)$	`Consumption`
fatigue	$f(t)$	`Fatigue`
operational cost at $t-1$	$o(t-1)$	`OperationalCost_1`
operational cost at $t-2$	$o(t-2)$	`OperationalCost_2`
operational cost at $t-3$	$o(t-3)$	`OperationalCost_3`
operational cost at $t-4$	$o(t-4)$	`OperationalCost_4`
operational cost at $t-5$	$o(t-5)$	`OperationalCost_5`
operational cost at $t-6$	$o(t-6)$	`OperationalCost_6`
operational cost at $t-7$	$o(t-7)$	`OperationalCost_7`
operational cost at $t-8$	$o(t-8)$	`OperationalCost_8`
operational cost at $t-9$	$o(t-9)$	`OperationalCost_9`
1^{st} latent variable of mis-calibration	m_1^l	`MisCalibrationDomain`
2^{nd} latent variable of mis-calibration	m_2^l	`MisCalibrationSystemResponse`
3^{rd} latent variable of mis-calibration	m_3^l	`MisCalibrationPhiIdx`
1^{st} latent variable fatigue	h_v	`FatigueLatentV`
2^{nd} latent variable fatigue	h_g	`FatigueLatentG`

Table 2: The preferred minimal markovian state. It fulfills the markov property with the minimum number of variables.

Modifier and Type	Method and Description
`DataVector`	`clone()` Returns a copy of the data vector.
`List <String>`	`getKeys()` Returns a list containing the data-vector dimension names.
`Double`	`getValue(String_key)` Returns the value for a given data-vector dimension.
`double[]`	`getValuesArray()` Returns a double[] array containing the values.
`void`	`setValue(String_key, double_value)` Sets the current value of a given data-vector dimension.

7.2 Interface Environment

This interface describes all relevant methods for implementing the dynamics of an environment.

text name or description	symbol	software name
set point	$p(t)$	`SetPoint`
velocity	$v(t)$	`Velocity`
effective velocity	v^e	`EffectiveVelocity`
gain	$g(t)$	`Gain`
effective gain	g^e	`EffectiveGain`
shift	$s(t)$	`Shift`
effective shift	$s^e(t)$	`EffectiveShift`
1^{st} latent variable of mis-calibration	m_1^l	`MisCalibrationDomain`
2^{nd} latent variable of mis-calibration	m_2^l	`MisCalibrationSystemResponse`
3^{rd} latent variable of mis-calibration	m_3^l	`MisCalibrationPhiIdx`
mis-calibration	$m(t)$	`MisCalibration`
modified operational cost	$\hat{c}$	`NoiseFreeConsumption`
consumption	$c(t)$	`Consumption`
fatigue	$f(t)$	`Fatigue`
current operational cost	$o(t)$	`OperationalCost_0`
operational cost at $t-1$	$o(t-1)$	`OperationalCost_1`
operational cost at $t-2$	$o(t-2)$	`OperationalCost_2`
operational cost at $t-3$	$o(t-3)$	`OperationalCost_3`
operational cost at $t-4$	$o(t-4)$	`OperationalCost_4`
operational cost at $t-5$	$o(t-5)$	`OperationalCost_5`
operational cost at $t-6$	$o(t-6)$	`OperationalCost_6`
operational cost at $t-7$	$o(t-7)$	`OperationalCost_7`
operational cost at $t-8$	$o(t-8)$	`OperationalCost_8`
operational cost at $t-9$	$o(t-9)$	`OperationalCost_9`
convoluted operational cost	$o^c(t)$	`OperationalCostConv`
modified operational cost	$\hat{c}(t)$	`ModifiedOperationalCost`
1^{st} latent variable fatigue	h_v	`FatigueLatentV`
2^{nd} latent variable fatigue	h_g	`FatigueLatentG`

Table 3: The extended state.

Modifier and Type	Method and Description
`void`	`addExternalDriver(ExternalDriver_extDriver)` This function adds an external driver to the environment, which affect/filter state variables during the call of step().
`DataVector`	`getInternalMarkovState()` Returns the internal Markovian state.
`double`	`getReward()` Returns the reward.
`DataVector`	`getState()` Returns the observable state.
`void`	`reset()` Function for resetting the environment.
`double`	`step(DataVector_action)` Performs an action within the environment and returns the reward.

7.3 Interface ExternalDriver

Abstract interface for attaching external drivers to the Environment, that affect/filter certain state dimensions (e.g. such as set point).

Modifier and Type	Method and Description
`void`	`filter(DataVector_state)` Applies "in-place" the external drivers to the given data vector.
`DataVector`	`getState()` Returns the current configuration.
`void`	`setConfiguration()` Sets the external driver configuration from within the given data vector.
`void`	`setSeed(long_seed)` Sets the random seed.

8 Example usage

8.1 Class ExampleMain

An example is implemented in `com.siemens.industrialsim.ExampleMain`.

Modifier and Type	Method and Description
`static void`	`main(String[]_args)` Run example simulation with random actions.

9 First results

First results with 10.000 data vectors as described in section 6 indicate, that the reward (named `RewardTotal` can be estimated from current and past values of velocity, gain, shift, and set point as inputs by a recurrent neural network with a mean relative absolute deviation (MRABD) of approximately 10%. Consumption c can be estimated with a MRABD of approximately 3.6%, and fatigue f with a MRABD of approximately 24% (for $f > 1$).

The average reward of the maximum entropy policy in the setting described in section 6 is -290.8 ± 0.6 with a standard deviation of 20.

First results with the policy gradient neural rewards regression (PGNRR) [4] and with an extension to continues actions of the neural fitted Q-iteration (NFQ) [5] lead to average rewards of roughly -270.

References

[1] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, 1998.

[2] S. Thrun. Is learning the n-th thing any easier than learning the first? In D. Touretzky and M. Mozer, editors, *Advances in Neural Information Processing Systems (NIPS) 8*, pages 640–646, Cambridge, MA, 1996. MIT Press.

[3] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, Oct 2010.

[4] Daniel Schneegaß, Steffen Udluft, and Thomas Martinetz. Improving optimality of neural rewards regression for data-efficient batch near-optimal policy identification. In Joaquim Marques de Sá, Luís A. Alexandre, Włodzisław Duch, and Danilo Mandic, editors, *Artificial Neural Networks – ICANN 2007: 17th International Conference, Porto, Portugal, September 9-13, 2007, Proceedings, Part I*, pages 109–118, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.

[5] Martin Riedmiller. Neural fitted Q-iteration - first experiences with a data efficient neural reinforcement learning method. In *Proceedings of the 16th European Conference on Machine Learning*, pages 317–328, 2005.